

Appl. No. 10/848,869
Amdt dated November 23, 2007

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application.

Listing of Claims:

1 (Currently amended): A method of using a plurality of row-identifier and value pairs to update rows in a table of a relational database, the method comprising:

repeatedly finding, and storing in a structure, a block-identifier of a block that contains a row of data identified by a row-identifier in at least a group of row-identifier and value pairs, by use of a database index prior to retrieval of the block;

wherein each value comprises the data in said row identified by said row-identifier;

performing a single access operation without context switching, to retrieve from a storage device and store in a buffer-cache, a number of blocks, said blocks being identified by a corresponding number of block-identifiers in the structure; and

repeatedly updating, in blocks in the buffer cache, each row identified in the group of row-identifier and value pairs, using a corresponding value in the row-identifier and value pairs.

2 (Currently amended): The method of Claim 1 further comprising:

sorting the block identifiers, prior to retrieval of the blocks by performing the vector-read single access operation.

3 (Original): The method of Claim 2 wherein:

SILICON VALLEY
PATENT GROUP LLP
15400 Cox Avenue
Suite 200
San Diego, CA 92120
(408) 378-7777
FAX (408) 378-7770

Appl. No. 10/848,869
Amdt dated November 23, 2007

the sorting is performed subsequent to storage of the block identifiers in the structure.

4 (Original): The method of Claim 1 further comprising:

subsequent to said finding and prior to said storing, checking if the block identifier has a duplicate already stored in the structure and if so then not storing the block identifier in the structure.

5 (Previously presented): The method of Claim 1 further comprising, prior to updating:

repeating said finding of block-identifiers for all row-identifiers in the group of row-identifier and value pairs.

6 (Original): The method of Claim 1 wherein:

the database index is a hash index and the table is organized in a hash cluster;
and

during said finding, a single directory is used to obtain the block identifier.

7 (Original): The method of Claim 1 wherein:

the database index is a B-tree index.

SILICON VALLEY
PATENT GROUP LLP

15505 Cus Avenue
Suite 220
San Jose, CA 95030
(408) 378-7777
FAX (408) 378-7770

Appl. No. 10/848,869
Amdt dated November 23, 2007

8 (Currently amended): The method of Claim 1 wherein:

said structure comprises an array; and

the array has a number of entries identical to the number of blocks that can be held in the buffer cache.

9 (Currently amended): The method of Claim 1 further comprising:

writing a plurality of logs, at least one log for each row identified in the group of row-identifier and value pairs and performing a write operation from said cache to said storage device when space is needed in said cache.

10 (Currently amended): The method of Claim 9 further comprising, during said write operation:

unpinning each block after updating all rows in said each block; and

flushing an unpinned block to disk only when another block needs space in the buffer cache occupied by the unpinned block.

11 (Currently amended): The method of Claim 1 wherein:

a plurality of file offsets are provided to the ~~vector read~~ single access operation, one offset for each block in the group.

Claim 12 (canceled).

SILICON VALLEY
PATENT GROUP LLP

11805 Cox Avenue
Suite 220
Saratoga, CA 95070
(408) 378-7777
FAX (408) 378-7770

Appl. No. 10/848,869

Amdt dated November 23, 2007

13 (Currently amended): A ~~computer-readable storage medium encoded with non-volatile media in which are stored~~ instructions to perform a method comprising:

repeatedly finding, and storing in a structure, a block-identifier of a block that contains a row identified by a row-identifier in at least a group of row-identifier and value pairs, by use of a database index of a relational database;

performing a vector read operation without context switching during said performing, to retrieve from a storage device and store in a buffer cache, a number of blocks, said blocks being identified by block-identifiers in the structure; and

repeatedly updating, in blocks in the buffer cache, each row identified in the group of row-identifier and value pairs, using a corresponding value in the row-identifier and value pairs.

14 (Currently amended): The ~~computer-readable storage medium non-volatile media~~ of Claim 13 being further encoded with said structure storing the block identifiers.

15 (Currently amended): A computer comprising a processor and a memory coupled to the processor, the memory being encoded with instructions to:

automatically use a database index to look up a block identifier of a block that contains a row identified by an identifier in a plurality of identifier and value pairs to be used to update a table in a relational database;

automatically store the block identifier in a structure in memory;

automatically repeat instructions to said automatically use and said automatically store, for all identifiers in at least a group of identifier and value pairs;

Appl. No. 10/848,869
Amdt dated November 23, 2007

automatically perform a vector read, to retrieve from a disk and store in a cache, each block in a group of blocks identified by block identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

automatically modify a row in a block stored in the cache, using a value in the plurality of identifier and value pairs; and

automatically repeat instructions to said automatically modify, with each row identified in the group of identifier and value pairs.

16 (Currently amended): An apparatus for using a plurality of identifier and value pairs to update a table of a database, each identifier in each pair identifying a row in the table, the apparatus comprising:

means for using a database index to look up a block identifier of a block that contains the row identified by an identifier in the plurality of identifier and value pairs;

means for storing the block identifier in a structure in memory;

means for repeating (using the database index to look up and storing the block identifier), for all identifiers in at least a group of identifier and value pairs;

means for performing a vector read without context switching, to retrieve from a disk and store in a cache, each block in a group of blocks identified by block identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

means for modifying a row in a block stored in the cache, using a value in the plurality of identifier and value pairs; and

SILICON VALLEY
PATENT GROUP LLP
18805 Cox Avenue
Suite 220
Saratoga, CA 95070
(408) 378-7777
FAX (408) 378-7770

Appl. No. 10/848,869
Amtdt dated November 23, 2007

means for repeating said modifying with each row identified in the group of identifier and value pairs.

17 (Currently amended): A method of using a plurality of row-identifier and value pairs to update a table of a database, each row-identifier in each pair identifying a row in the table, the method comprising:

finding a block-identifier of a block that contains the row identified by a row-identifier in a row-identifier and value pair, by use of a database index;

storing the block-identifier in a structure;

repeating (finding the block-identifier and storing the block-identifier), for all row-identifiers in at least a group of row-identifier and value pairs;

performing a vector read operation without context switching during said performing, to retrieve from a storage device and store in a ~~buffer~~ cache, each block in a group of blocks identified by block-identifiers stored in said structure, wherein the group of blocks are all stored in the cache during execution of a single function call;

updating the row in the block in the cache, using the value in the row-identifier and value pair; and

repeating said updating with each row identified in the group of row-identifier and value pairs.

18 (New): The non-volatile media of Claim 13 being comprised in at least one of a disk, a chip and a cartridge.

SILICON VALLEY
PATENT GROUP LLP
15395 Cope Avenue
Suite 220
San Jose, CA 95070
(408) 378-7777
FAX (408) 378-7770

Appl. No. 10/848,869
Amdt dated November 23, 2007

19 (New): The method of Claim 2 wherein:

the blocks are sorted during said sorting based on adjacency such that during performance of said single access operation, block identifiers of blocks physically adjacent to one another at a periphery of a disk in the storage device are presented at one time to the storage device and identifiers of blocks that are physically adjacent to one another and located closer to a center of the disk are presented at another time.

20 (New): The computer of Claim 15 wherein:

the blocks are sorted during said single function call based on adjacency such that block identifiers of blocks physically adjacent to one another at a periphery of said disk are presented at one time to a disk drive comprising said disk and identifiers of blocks that are physically adjacent to one another and located closer to a center of said disk are presented at another time.

SILICON VALLEY
PATENT GROUP LLP

18805 Cox Avenue
Suite 220
San Jose, CA 95070
(408) 378-7771
FAX (408) 378-7770